

SANDIA REPORT

SAND2007-5984

Unlimited Release

Printed September 2007

Thyra Linear Operators and Vectors

Overview of Interfaces and Support Software for the Development and Interoperability of Abstract Numerical Algorithms

Roscoe A. Bartlett

Optimization and Uncertainty Quantification

Prepared by

Sandia National Laboratories

Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia is a multiprogram laboratory operated by Sandia Corporation,
a Lockheed Martin Company, for the United States Department of Energy's
National Nuclear Security Administration under Contract DE-AC04-94-AL85000.

Approved for public release; further dissemination unlimited.



Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: <http://www.osti.gov/bridge>

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd
Springfield, VA 22161

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.fedworld.gov
Online ordering: <http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online>



Thyra Linear Operators and Vectors

Overview of Interfaces and Support Software for the Development and Interoperability of Abstract Numerical Algorithms

Roscoe A. Bartlett
Optimization/Uncertainty Estim

Abstract

Engineering and scientific applications are becoming increasingly modular, utilizing publicly defined interfaces to integrate third party tools and libraries for services such as mesh generation, data partitioning, equation solvers and optimization. As a result, it is important to understand and model the interactions between these various modules, and to develop good abstract interfaces between them. One category of modules that is becoming increasingly important is abstract numerical algorithms (ANAs). ANAs such as linear and nonlinear equation solvers, methods for stability and bifurcation analysis, uncertainty quantification methods and nonlinear programming solvers for optimization are typically mathematically sophisticated but have surprisingly little essential dependence on the details of what computer system is being used or how matrices and vectors are stored and computed. As a result, using abstract interface capabilities in languages such as C++, we can implement ANA software that it will work, unchanged, with a variety of applications and linear algebra libraries.

In this paper, we provide an overview of the Thyra effort which at its most basic level defines fundamental abstract linear operator and vector interfaces. These linear operator/vector interfaces provide the basic functionality and interoperability for a broad range of ANAs. Many other higher-level abstractions are built on top of the Thyra operator/vector interfaces. The Trilinos package `thyra` defines these different sets of C++ interfaces and provides optional support software.

Acknowledgment

The authors would like to thank everyone on the Trilinos team for their support in the Thyra effort.

The format of this report is based on information found in [15].

Contents

1	Introduction	7
2	Classification of linear algebra and other interfaces	8
3	Some Basic Requirements for Thyra	11
4	Mathematical Foundation for Fundamental Thyra ANA Operator/Vector Interfaces	12
5	Overview of Fundamental Thyra ANA Operator/Vector Interfaces	14
6	Summary	17
	References	18

Figures

- | | | |
|---|---|----|
| 1 | UML [6] class diagram : Interfaces between abstract numerical algorithm (ANA), linear algebra library (LAL), and application (APP) software. | 10 |
| 2 | UML class diagram : The fundamental Thyra ANA operator/vector interfaces | 16 |

1 Introduction

One area of steady improvement in large-scale engineering and scientific applications is the increased modularity of application design and development. Specification of publicly-defined interfaces, combined with the use of third-party software to satisfy critical technology needs in areas such as mesh generation, data partitioning and solution methods have been generally positive developments in application design. While the use of third party software introduces dependencies from the application developer's perspective, it also gives the application access to the latest technology in these areas, amortizes library and tool development across multiple applications and, if properly designed, gives the application easy access to more than one option for each critical technology area, e.g., access to multiple linear solver packages.

One category of modules that is becoming increasingly important is abstract numerical algorithms (ANAs). ANAs such as linear and nonlinear equation solvers, methods for stability and bifurcation analysis, transient solvers, uncertainty quantification methods, and nonlinear programming solvers for optimization are typically mathematically sophisticated but have surprisingly little essential dependence on the details of what computer system is being used or how matrices and vectors are stored and computed. Thus, by using abstract interface capabilities in languages such as C++, we can implement ANA software such that it will work, unchanged, with a variety of applications and linear algebra libraries.

Here we describe a set of abstract operator/vector interfaces that allows the specification of ANAs from basic Krylov linear equation solvers all the way up to interior-point methods for optimization. At the core, we define a set of basic operator/vector interfaces that form the the foundation for (i) ANA development, (ii) the integration of an ANA into an application (APP) and (iii) providing services to the ANA from a linear algebra library (LAL). By agreeing on a simple minimal common interface layer such as the Fundamental Thyra ANA Operator/Vector Interfaces described here, we eliminate the many-to-many dependency problem of ANA/APP interfaces.

It is difficult to describe a set of linear algebra interfaces outside of the context of some class of numerical problems. For this purpose, we will consider numerical algorithms where it is possible to implement all of the required operations exclusively through well defined interfaces to vectors, vector spaces, and linear operators and higher level abstractions built on these. The fundamental Thyra operator/vector interfaces described here are the common denominator of all abstract numerical algorithms.

We assume that the reader has a basic understanding of vector reduction/transformation operators (RTOp) [3], is comfortable with object-orientation [11] and C++, and knows how to read basic Unified Modeling Language (UML) [10] class diagrams. We also assume that the reader has some background in large-scale numerics and will therefore be able to appreciate the challenges that are addressed by Thyra.

Note that the online documentation for Thyra at

<http://trilinos.sandia.gov/packages/thyra>

should be the definitive information source for Thyra. This document only tries to provide an overview of Thyra and explain the philosophy behind it.

2 Classification of linear algebra and other interfaces

Although we will discuss APPs, ANAs and LALs in detail later in this section, we want to briefly introduce these terms here to make them clear. Also, although there are certainly other types of modules in a large-scale scientific/engineering application, we only focus on these three since they are the ones more directly related to ANAs.

- **Application (APP):** The modules of an application that are not ANA or LAL modules. Typically this includes the code that is unique to the application itself such as the code that formulates and generates the discrete problem to be solved. In general it would also include other third-party software that is not an ANA or LAL module.
- **Abstract Numerical Algorithm (ANA):** Software that drives a solution process, e.g., an iterative linear or nonlinear solver. This type of package provides solutions to and requires services from the APP, and utilizes services from one or more LALs. It can usually be written so that it does not depend on the details of the computer platform, or the details of how the APP and LALs are implemented, so that an ANA can be used across many APPs and with many LALs.
- **Linear Algebra Library (LAL):** Software that provides the ability to construct concrete linear algebra objects such as matrices and vectors. A LAL can also be a specific linear solver or preconditioner.

An important focus of this paper is to clearly identify the interactions between APPs, ANAs and LALs for the purposes of defining the Thyra interfaces and to differentiate the Thyra interfaces from other interfacing efforts.

The requirements for the linear algebra objects as imposed by an ANA are very different from the requirements imposed by an APP code. In order to differentiate the various types of interfaces and the requirements associated with each, consider Figure 1. This figure shows the three major categories of software modules that make up a complete numerical application. The first category is application (APP) software in which the underlying data is defined for the problem. This could be something as simple as the right-hand-side and matrix coefficients of a single linear system or as complex as a finite-element method for a 3-D nonlinear PDE-constrained optimization problem. The second category is linear algebra library (LAL) software that implements basic linear algebra operations [9, 1, 5, 13, 2, 12]. These types of software include primarily matrix-vector multiplication, the creation of a preconditioner (e.g. ILU), and may even include several different types of direct linear solvers. The third category is ANA software that drives the main solution process and includes such algorithms as iterative methods for linear and nonlinear systems; explicit and implicit methods for ODEs and DAEs; and nonlinear programming (NLP) solvers [16]. There are many example software packages [2, 13, 12, 7, 4] that contain ANA software.

The types of ANAs described here only require operations like matrix-vector multiplication, linear solves and certain types of vector reduction and transformation operations. All of these operations can be performed with only a very abstract view of vectors, vector spaces and linear operators.

An application code, however, has the responsibility of populating vector and matrix objects and requires the passing of explicit function and gradient value entries, sometimes in a distributed

memory parallel environment. This is the purpose of an APP/LAL interface. This involves a very different set of requirements than those described above for the ANA/APP and ANA/LAL interfaces. Examples of APP/LAL interfaces include the FEI [8] and much of the ESI [14].

Figure 1 also shows a set of LAL/LAL interfaces that allows linear algebra objects from one LAL to collaborate with the objects from another LAL. These interfaces are very similar to the APP/LAL interfaces and the requirements for this type of interface is also not addressed by Thyra. The ESI [14] contains examples of LAL/LAL interfaces.

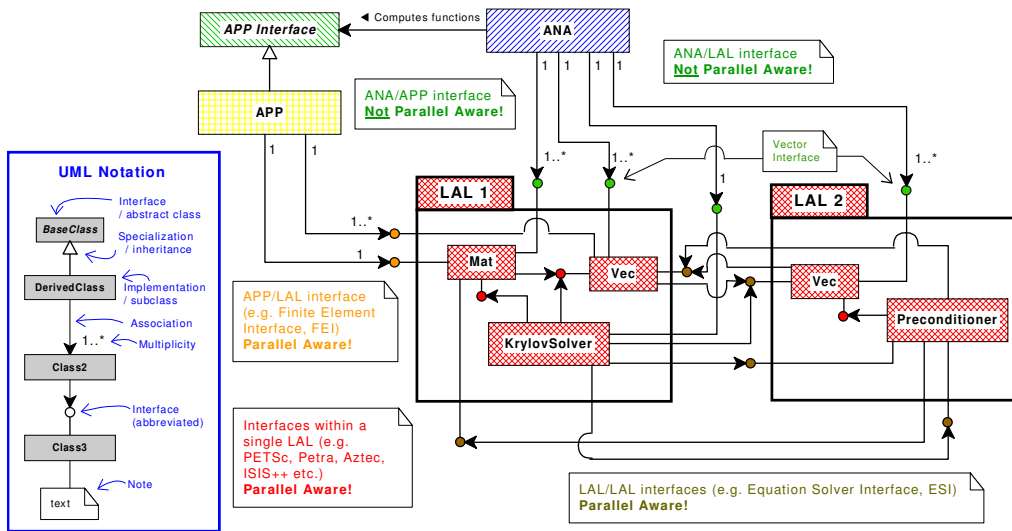


Figure 1. UML [6] class diagram : Interfaces between abstract numerical algorithm (ANA), linear algebra library (LAL), and application (APP) software.

3 Some Basic Requirements for Thyra

Before describing the C++ interfaces for Thyra, some basic requirements are stated.

1. Thyra interfaces should be portable to all the ASC [17] platforms where ASC applications might run and to other important platforms.
2. Thyra interfaces should provide for stable and accurate numerical computations at a fundamental level.
3. Thyra should provide a minimal, but complete, interface that addresses all the basic efficiency needs (in both speed and storage) which will result in near-optimal implementations of all of the objects and all of the above mentioned ANA algorithms that use these objects.
4. Maximally general ANAs developed with Thyra should be able to transparently utilize different types of computing environments such as SPMD¹, client/server², out-of-core³, and any combination of these configurations.
5. The work required to implement adapter subclasses (see the “Adapter” pattern in [11]) for and with Thyra should be minimal and straightforward for all of the existing related linear algebra and ANA interfaces. This requirement is facilitated by the fact that the Thyra interfaces are minimal.

A hand-coded program (e.g. using Fortran 77 and MPI) should not provide any significant gains in performance in any of the above categories in any computing environment or configuration. A hand-coded algorithm in Fortran 77 with MPI should not be able to provide significant improvements in storage requirements, computational speed, or numerical stability. There are many numerical algorithms that can not be considered to be “abstract” (e.g. Gaussian Elimination) and therefore Thyra and like abstract interfaces should not be used for such algorithms. However, drawly the line between an ANA and a non-ANA can be quite fuzzy in practice.

¹Single Program Multiple Data (SPMD): A single program running in a distributed-memory environment on multiple parallel processors

²Client/Server: The ANA runs in a process on a client computer and the APP and LAL run in processors on a server

³Out-of-core: The data for the problem is stored on disk and is read from and written to back disk as needed

4 Mathematical Foundation for Fundamental Thyra ANA Operator/Vector Interfaces

Before describing the *Thyra Fundamental ANA Operator/Vector Interfaces* that form the foundation for all of the Thyra ANA interfaces in more detail, we must first clearly define the mathematical foundation for the abstractions of vectors, vector spaces, and linear operators.

All Thyra vectors belong to a vector space and are representable as an array of scalar coefficients $\tilde{x} \in \mathbf{R}^n$ and a (non-unique) basis $E \in \mathbf{R}^{n \times n}$ of the form

$$x = E\tilde{x}. \quad (1)$$

For a given Thyra vector space, the basis representation $E \in \mathbf{R}^{n \times n}$ is non-unique but the symmetric positive definite scalar (inner) product matrix

$$Q = E^H E \quad (2)$$

where $Q \in \mathbf{R}^{n \times n}$ is unique. Given $Q \in \mathbf{R}^{n \times n}$, the scalar product is defined as

$$x^H y = \tilde{x}^H Q \tilde{y} = \langle \tilde{x}, \tilde{y} \rangle. \quad (3)$$

For $Q \in \mathbf{R}^{n \times n}$ to be full rank, $E \in \mathbf{R}^{n \times n}$ must be full rank. Here we define a Euclidean vector space as one where $E = I \in \mathbf{R}^{n \times n}$ and $Q = E^H E = I \in \mathbf{R}^{n \times n}$.

In Thyra, vector spaces and vectors are abstracted using the C++ base interface classes `Thyra::VectorSpaceBase` and `Thyra::VectorBase` respectively. Vectors are created from a vector space using the “Abstract Factory” design pattern using the nonmember function `createMember(vecSpc)`.

In addition to vectors and vector spaces, Thyra also defines linear operators which linearly map vectors from one vector space to another. The definition of a linear operator is strongly influenced by the definition of the scalar product associated with the vector spaces and whether the mapping is between Euclidean vectors $y = Ax$ or between coefficient vectors $\tilde{y} = \tilde{A}\tilde{x}$. In Thyra, vectors are a specialization of linear operators and therefore every vector is also a linear operator. Therefore, when one writes

$$z = x^H y \quad (4)$$

then x^H must be interpreted to be the adjoint linear operator of x .

Given this notation, the vector x would be considered to live in a Euclidean vector space $x \in \mathbf{R}^n$ while the vector \tilde{x} would be considered to live in the non-Euclidean vector space $\tilde{x} \in X$ which is defined by the scalar product matrix $Q \in \mathbf{R}^{n \times n}$. Therefore, it is actually not clear whether an abstract vector object represents the Euclidean vector $x \in \mathbf{R}^n$ which just happens to be stored as a set of coefficients $\tilde{x} \in X$ or if it represents the coefficient vectors $\tilde{x} \in X$ themselves. This ambiguity

of the interpretation of vectors and linear operators also results from the ambiguity of interpretation of the vectors in that they can be thought of as the Euclidean linear operators that only update the coefficient vectors or as linear operators that act directly on vector coefficients.

The Fundamental Thyra ANA Operator/Vector Interfaces require that every Thyra vector be represented as a finite-dimensional set of scalar coefficients and that the scalar product $\langle \tilde{x}, \tilde{y} \rangle$ be equivalent to the two-sided application some finite-dimensional symmetric positive definite matrix $Q \in \mathbf{R}^{n \times n}$ such that $\langle \tilde{x}, \tilde{y} \rangle = \tilde{x}^H Q \tilde{y}$. The Thyra interfaces do not try to pretend that its vectors are infinite dimensional or that they admit more general implementations as allowed by infinite dimensional Hilbert spaces. While every Thyra vector must be stored as a set of scalar coefficients the interfaces make no assumptions whatsoever about where or how those coefficients are stored. A fully general ANA can make no assumptions about how vectors are stored or laid out in memory, only that those coefficients do exist and that the coefficients can be exposed to reduction and transformation operators (see [3]).

While accessing the elements of a vector is ill advised in a general ANA, the Thyra vector interface supports acquiring direct views of any range of vector coefficient data (see `Thyra::VectorBase::acquireDetachedView(...)`).

Again, thyra does not try to hide the fact that a vector is simply a set of scalar coefficients associated with some basis. Thyra does not even really try to keep the client from accessing the actual coefficients of the vectors since it can access them with an RTOp object if desired. What Thyra does do is to abstract where the vector coefficients live and what native data structure is used to hold the coefficients. None of vector coefficients may even be directly be head in main memory in the process where the ANA is running but there is always a (perhaps very inefficient) mechanism to get a view of them. ANAs that want to be maximally general and efficient should not try to directly access the vector elements explicitly and many ANAs do not need to. However, there are perfectly reasonable ANAs that do need to access the explicit vector coefficients for vectors from certain vector spaces (such as vectors in the design space in some reduced-space optimization algorithms) and Thyra provides efficient and yet 100% general ways to access these coefficients. Every Thyra vector space also has a finite dimension that is accessible as a property of the vector space and in integer value.

While Thyra requires vectors to be finite dimensional and the vector coefficients must be accessible (if to none other than to RTOps), it allows complete freedom in the implementation of scalar products of general linear operators that map vectors from one vector space to another.

5 Overview of Fundamental Thyra ANA Operator/Vector Interfaces

The Fundamental Thyra ANA Operator/Vector Interfaces are shown in Figure 2. The key abstractions include vectors, vector spaces, and linear operators. All of the interfaces are templated on the `Scalar` type (the UML notation for templated classes is not used in the figure for the sake of improving readability).

Vector space is the foundation for all other abstractions. Vector spaces are abstracted through the *VectorSpaceBase* interface. A *VectorSpaceBase* object acts primarily as an “Abstract Factory” [11] that creates vector objects (which are the “products” in the “Abstract Factory” design pattern).

Vectors are abstracted through the *VectorBase* interface. The *VectorBase* interface is very minimal and really only defines one nontrivial function *applyOp(...)*. The *applyOp(...)* function accepts user-defined (i.e. ANA-defined) reduction/transformation operator (RTOp) objects through the templated RTOp C++ interface *RTOpPack: :RTOpT*. An ever increasing set of concrete implementations of RTOps is provided along with wrapper convenience functions in the ANA support code collection. The set of operations is also easily extensible. Every *VectorBase* object provides access to its *VectorSpaceBase* (that was used to create the *VectorBase* object) through the function *space()*.

The *VectorSpaceBase* interface also provides the ability to create *MultiVectorBase* objects through the *createMembers(numMembers)* function. A *MultiVectorBase* is a tall thin dense matrix where each column in the matrix is a *VectorBase* object which is accessible through the *col(...)* function. *MultiVectorBases* are needed for near-optimal processor cache performance (in serial and parallel programs) and to minimize the number of global communications in a distributed parallel environment. The *MultiVectorBase* interface is useful in many different types ANAs such as block Krylov methods. The interface class *VectorBase* is derived from *MultiVectorBase* so that every *VectorBase* is a *MultiVectorBase*. This simplifies the development of ANAs in that any ANA that can handle *MultiVectorBase* objects should automatically be able to accept *VectorBase* objects as well.

VectorSpaceBase declares a virtual function called *scalarProd(x,y)* which computes the scalar product $\langle x, y \rangle$ for the vector space. There is also a *MultiVectorBase* version *VectorSpaceBase::scalarProds(...)* (not shown in the figure) that computes the scalar products of each set of column vectors in two multi-vectors. Finally, *VectorSpaceBase* also includes the ability to determine the compatibility of vectors from different vector spaces through the function *isCompatible(vecSpc)*. This is useful primarily for error checking and debugging.

Another important type of linear algebra abstraction is a linear operator which is represented by the interface class *LinearOpBase*. The *LinearOpBase* interface is used to represent quantities such as a Jacobian matrix. A *LinearOpBase* object defines a linear mapping from vectors in one vector space (called the domain) to vectors in another vector space (called the range). Every *LinearOpBase* object provides access to these vector spaces through the functions *domain()*

and `range()`. The exact form of this mapping, as implemented by the function `apply(...)`, is

$$Y = \alpha M X + \beta Y \quad (5)$$

where M is a *LinearOpBase* object; X and Y are *MultiVectorBase* objects; and α and β are Scalar objects.

A *LinearOpBase* object can also, optionally, support the transpose (or adjoint) operation:

$$Y = \alpha op(M) X + \beta Y \quad (6)$$

through the function `applyTranspose(...)`, where M is a *LinearOpBase* object; $op(M)$ is M^T or M^H (as determined by the *conj* argument); X and Y are *MultiVectorBase* objects; and α and β are Scalar objects. The function `applyTranspose(EConj)` will return false if a particular form of the transpose is not supported.

If the adjoint is supported, then it must satisfy the adjoint property. Specifically, for any two vectors $w \in \mathcal{D}$ (domain space) and $u \in \mathcal{R}$ (range space), the adjoint operation must obey the adjoint property

$$\langle u, Av \rangle_{\mathcal{R}} = \langle A^H u, v \rangle_{\mathcal{D}}.$$

Another important part of this design is the fact that *MultiVectorBase* derives from *LinearOpBase* and therefore every multi-vector object is also a linear operator. This is an elegant way to support the notions of block inner products and block updates.

A block inner product is specified as

$$Z = Y^H X$$

where Y , X and Z are all multi-vectors. Note that since Y is a linear operator then $Y^H X$ is not simply the block dot product involving the coefficients but instead must be consistent with the scalar product for the range of Y .

A block update takes the form

$$Z = \alpha Y X + \beta Z$$

where Y , X and Z are all multi-vectors and α and β are scalars.

Also note that since *MultiVectorBase* derives from *LinearOpBase* and *VectorBase* derives from *MultiVectorBase*, therefore, every vector object is also a linear operator. While this may not be a terribly useful feature it does mean that one must interpret $y^H x$ to be the same as $\langle y, x \rangle$ and not just the dot product when the space is non-euclidean.

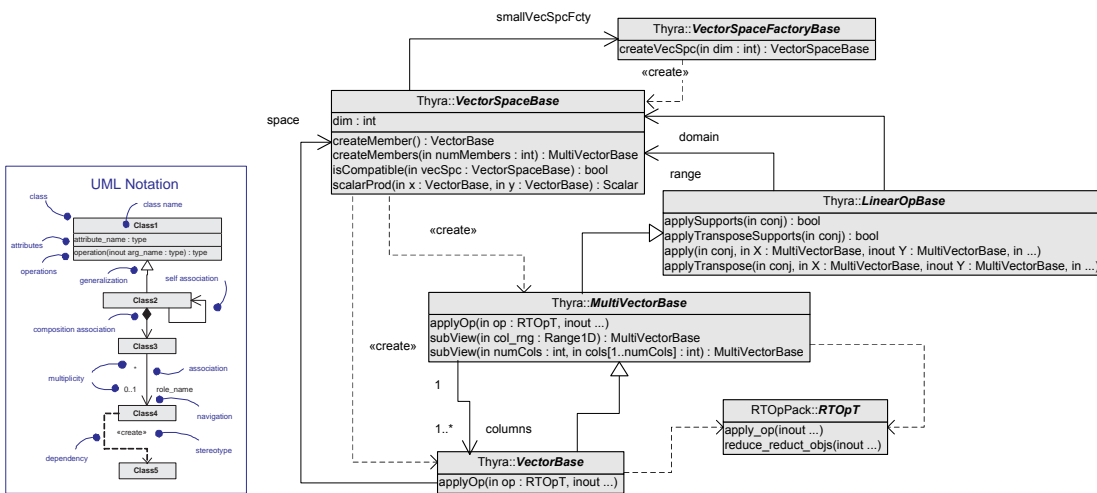


Figure 2. UML class diagram : The fundamental Thyra ANA operator/vector interfaces

6 Summary

The Fundamental Thyra Operator/Vector Interfaces provide the intersection of all of the functionality required of linear operators and vectors by a variety of abstract numerical algorithms ranging from iterative linear solvers all the way up to optimizers. By adopting Thyra as a standard interface layer, interoperability between applications, linear algebra libraries, and abstract numerical algorithms in advanced scientific computing environments becomes automatic to a large extent.

References

- [1] E. Anderson et al. *LAPACK User's Guide*. SIAM, 1995.
- [2] S. Balay, W. D. Gropp, L.C. McInnes, and B.F. Smith. PETSc 2.0. <http://www.mcs.anl.gov/petsc>.
- [3] R. A. Bartlett, B. G. van Bloeman Waanders, and M. A. Heroux. Vector reduction/transformation operators for linear algebra interfaces to efficiently develop complex abstract numerical algorithms independently of data mapping, 2003. Submitted to *ACM TOMS*.
- [4] Steve Benson, Lois Curfman McInnes, and Jorge Moré. TAO : Toolkit for advanced optimization (web page).
- [5] L. S. Blackford, J. Choi, A. Cleary, E.D. 'Azevedo, J. Demmel, I. Dhillon, J. Dongarra, S. Hammarling, G. Henry, A. Petitet, K. Stanley, D. Walder, and R.C. Whaley. *ScalLAPACK User's Guide*. SIAM, Philadelphia, PA, 1997.
- [6] G. Booch, J. Rumbaugh, and I. Jacobson. *The Unified Modeling Language User Guide*. Addison-Wesley, 1999.
- [7] George D. Byrne and Allan C. Hindmarsh. PVODE, an ODE solver for parallel computers. *Int. J. High Perf. Comput. Applic.*, 13:354–365, 1999.
- [8] Robert L. Clay, Kyran D. Mish, Ivan J. Otero, Lee M. Taylor, and Alan B. Williams. An annotated reference guide to the finite-element interface (FEI) specification : Version 1.0. Technical Report SAND99-8229, Sandia National Laboratories, 1999.
- [9] J. Demmel. *Applied Numerical Linear Algebra*. SIAM, 1997.
- [10] M. Fowler and K. Scott. *UML Distilled, second edition*. Addison-Wesley, 2000.
- [11] E. Gamma et al. *Design Patterns: Elements fo Reusable Object-Oriented Software*. Addison-Wesley, 1995.
- [12] Mike A. Heroux, Teri Barth, David Day, Rob Hoekstra, Rich Lehoucq, Kevin Long, Roger Pawlowski, Ray Tuminaro, and Alan Williams. Trilinos : object-oriented, high-performance parallel solver libraries for the solution of large-scale complex multi-physics engineering and scientific applications. <http://software.sandia.gov/Trilinos>.
- [13] S. A. Hutchinson, L. V. Prevost, J.N. Shadid, C. Tong, and R.S. Tuminaro. Aztec user's guide: Version 2.0. Technical Report ANL-95/11–Revision 2.0.22, Sandia National Laboratories, 1998.
- [14] Sandia National Labs. ESI: Equation solver interface. <http://z.ca.sandia.gov/esi>, 2001.
- [15] Tamara K. Locke. Guide to preparing SAND reports. Technical report SAND98-0730, Sandia National Laboratories, Albuquerque, New Mexico 87185 and Livermore, California 94550, May 1998.
- [16] J. Nocedal and S. Wright. *Numerical Optimization*. Springer, New York, 1999.

- [17] Department of Energy. ASCI: Advanced simulation and computing initiative.
<http://www.sandia.gov/ASCI>.

DISTRIBUTION:

- 2 MS 9018 Central Technical Files, 8944
- 2 MS 0899 Technical Library, 4536

